

WHAT IS CLAIMED IS:

1. A speech recognition method comprising:
 - obtaining a set of acoustic observations;
 - obtaining a list of target speech element sequences each containing at least one speech element;
 - for each target speech element sequence obtaining a forward sequence extension model and a backward sequence extension model;
 - spotting at least one spotted target speech element sequence by matching the sequence of speech element models against the set of acoustic observations;
 - obtaining from the set of acoustic observations the set of acoustic observations preceding the said at least one spotted target speech element sequence and the set of acoustic observations following the said at least one spotted target speech element sequence;

obtaining at least one hypothesis of a longer speech element sequence containing the said at least one spotted speech element sequence as a proper subsequence in which said at least one longer speech element sequence is consistent with at least one of said forward sequence extension model and said backward sequence extension model for said at least one spotted speech element sequence; and

evaluating said at least one hypothesis of a longer speech element sequence based on the degree of acoustic match between said longer speech element sequence and at least one of said set of acoustic observations preceding the said at least one spotted target speech element sequence and the set of acoustic observations following the said at least one spotted target speech element sequence.

2. A speech recognition method as in claim 1, further comprising:

spotting a plurality of spotted target speech element sequences in the set of acoustic observations;

determining, for each spotted speech element sequence and each hypothesized longer speech element sequence, the set of acoustic observations that correspond to the speech interval for said speech element sequence;

detecting when the set of acoustic observations for a first speech element sequence and the set of acoustic observations for a second speech element sequence correspond to adjacent speech intervals; and

creating a combined speech element sequence by concatenating said first speech element sequence and said second speech element sequence.

3. A speech recognition method as in claim 2, further comprising:

obtaining from the set of acoustic observations the set of acoustic observations preceding the said at least one combined speech element sequence and the set of acoustic observations following the said at least one combined speech element sequence;

obtaining at least one hypothesis of a longer speech element sequence containing the said at least one combined speech element sequence as a proper subsequence in which said at least one longer speech element sequence is consistent with at least one of said forward sequence extension model of the spotted target speech element sequence contained in said second speech element sequence and said backward sequence extension model for the spotted target speech element sequence contained in said first speech element sequence; and

evaluating said at least one hypothesis of a longer speech element sequence based on the degree of acoustic match between said longer speech element sequence and at least one of said set of acoustic observations preceding the said at least one combined speech element sequence and the set of acoustic observations following the said at least one combined speech element sequence.

4. A speech recognition method as in claim 3, further comprising:

repeating said processes of obtaining at least one hypothesis of a longer speech element sequence, and said evaluating said at least one hypothesis, and said determining of said sets of corresponding acoustic observations, until there is at least one pair of a first speech element sequence and a second element sequence for which it is detected that said first speech element sequence and said second element sequence correspond to adjacent speech intervals;

creating said combined speech element sequence; and

repeating said processes of obtaining and evaluating said longer speech element sequences and of creating said combined speech element sequences until there is at least one hypothesized speech element sequence that corresponds to the complete set of acoustic observations.

5. A speech recognition method as in claim 1, further comprising:

obtaining a grammar of the allowed speech element sequences;

for each allowed target speech element sequence,
determining from the grammar the set of predecessor speech
element sequences that may precede said target speech element
sequence as adjacent subsequences in an allowed speech
element sequence;

creating a backward sequence extension model for said
target speech element sequence from said set of predecessor
speech element sequences;

for each target speech element sequence, determining
from the grammar the set of successor speech element
sequences that may follow said target speech element sequence
as adjacent subsequences in an allowed speech element
sequence; and

creating a forward sequence extension model for said
target speech element sequence from said set of successor
speech element sequences.

6. A speech recognition method as in claim 5, wherein said speech element sequences are word sequences and said grammar is a grammar of allowed word sequences.

7. A speech recognition method as in claim 1, wherein each target speech element sequences is a target phoneme sequence, and wherein the method further comprising:

obtaining a vocabulary list of speech elements each of which is a sequence of phonemes;

for each target phoneme sequence, determining from said vocabulary list the set of predecessor phoneme sequences that may precede said target phoneme sequences as an adjacent phoneme subsequence in the set of phoneme sequences in said vocabulary list;

creating a backward sequence extension model for said target phoneme sequence from said set of predecessor phoneme sequences; and

for each target phoneme sequence, determining from said vocabulary list the set of successor phoneme sequences that may follow said target phoneme sequence as an adjacent phoneme subsequence in the set of phoneme sequences in said vocabulary list.

8. A speech recognition method as in claim 1, wherein the set of acoustic observations is a sequence, and wherein the method further comprising:

performing a sequential speech recognition search substantially simultaneously with said spotting of at least one target speech element sequence; and

using said spotting of at least one speech element sequence to enhance said sequential speech recognition search.

9. A speech recognition method as in claim 8, wherein said sequential speech recognition search is a priority queue search.

10. A speech recognition method as in claim 8, wherein said sequential speech recognition search is a frame synchronous beam search.

11. A speech recognition system, comprising:

means for obtaining a list of target speech element sequences from a set of acoustic observations, each said target speech element sequence containing at least one speech element;

means for obtaining, for each said target speech element sequence, a forward sequence extension model and a backward sequence extension model;

means for spotting at least one spotted target speech element sequence by matching the sequence of speech element models against the set of acoustic observations;

means for obtaining, from the set of acoustic observations, the set of acoustic observations preceding the said at least one spotted target speech element sequence and

the set of acoustic observations following the said at least one spotted target speech element sequence;

means for obtaining at least one hypothesis of a longer speech element sequence containing the said at least one spotted speech element sequence as a proper subsequence in which said at least one longer speech element sequence is consistent with at least one of said forward sequence extension model and said backward sequence extension model for said at least one spotted speech element sequence; and

means for evaluating said at least one hypothesis of a longer speech element sequence based on the degree of acoustic match between said longer speech element sequence and at least one of said set of acoustic observations preceding the said at least one spotted target speech element sequence and the set of acoustic observations following the said at least one spotted target speech element sequence.

12. A speech recognition system as in claim 11, further comprising:

means for spotting a plurality of spotted target speech element sequences in the set of acoustic observations;

means for determining, for each spotted speech element sequence and each hypothesized longer speech element sequence, the set of acoustic observations that correspond to the speech interval for said speech element sequence;

means for detecting when the set of acoustic observations for a first speech element sequence and the set of acoustic observations for a second speech element sequence correspond to adjacent speech intervals; and

means for creating a combined speech element sequence by concatenating said first speech element sequence and said second speech element sequence.

13. A speech recognition system as in claim 12, further comprising:

means for obtaining from the set of acoustic observations the set of acoustic observations preceding the said

at least one combined speech element sequence and the set of acoustic observations following the said at least one combined speech element sequence;

means for obtaining at least one hypothesis of a longer speech element sequence containing the said at least one combined speech element sequence as a proper subsequence in which said at least one longer speech element sequence is consistent with at least one of said forward sequence extension model of the spotted target speech element sequence contained in said second speech element sequence and said backward sequence extension model for the spotted target speech element sequence contained in said first speech element sequence; and

means for evaluating said at least one hypothesis of a longer speech element sequence based on the degree of acoustic match between said longer speech element sequence and at least one of said set of acoustic observations preceding the said at least one combined speech element sequence and

the set of acoustic observations following the said at least one combined speech element sequence.

14. A speech recognition system as in claim 13, further comprising:

means for repeating said processes of obtaining at least one hypothesis of a longer speech element sequence, and said evaluating said at least one hypothesis, and said determining of said sets of corresponding acoustic observations, until there is at least one pair of a first speech element sequence and a second element sequence for which it is detected that said first speech element sequence and said second element sequence correspond to adjacent speech intervals;

means for creating said combined speech element sequence; and

means for repeating said processes of obtaining and evaluating said longer speech element sequences and of creating said combined speech element sequences until there is

at least one hypothesized speech element sequence that corresponds to the complete set of acoustic observations.

15. A speech recognition system as in claim 11, further comprising:

means for obtaining a grammar of the allowed speech element sequences;

means for determining, from the grammar for each allowed target speech element sequence, the set of predecessor speech element sequences that may precede said target speech element sequence as adjacent subsequences in an allowed speech element sequence;

means for creating a backward sequence extension model for said target speech element sequence from said set of predecessor speech element sequences;

means for determining from the grammar, for each target speech element sequence, the set of successor speech element sequences that may follow said target speech element

sequence as adjacent subsequences in an allowed speech element sequence; and

means for creating a forward sequence extension model for said target speech element sequence from said set of successor speech element sequences.

16. A speech recognition system as in claim 15, wherein said speech element sequences are word sequences and said grammar is a grammar of allowed word sequences.

17. A speech recognition system as in claim 11, wherein each target speech element sequences is a target phoneme sequence, and wherein the system further comprising:

means for obtaining a vocabulary list of speech elements each of which is a sequence of phonemes;

means for determining from the vocabulary list, for each target phoneme sequence, the set of predecessor phoneme sequences that may precede said target phoneme sequences as

an adjacent phoneme subsequence in the set of phoneme sequences in said vocabulary list;

means for creating a backward sequence extension model for said target phoneme sequence from said set of predecessor phoneme sequences; and

means for determining from the vocabulary list, for each target phoneme sequence, the set of successor phoneme sequences that may follow said target phoneme sequence as an adjacent phoneme subsequence in the set of phoneme sequences in said vocabulary list.

18. A speech recognition system as in claim 11, wherein the set of acoustic observations is a sequence, and wherein the system further comprising:

means for performing a sequential speech recognition search substantially simultaneously with said spotting of at least one target speech element sequence; and

means for using said spotting of at least one speech element sequence to enhance said sequential speech recognition search.

19. A speech recognition system as in claim 18, wherein said sequential speech recognition search is a priority queue search.

20. A speech recognition system as in claim 18, wherein said sequential speech recognition search is a frame synchronous beam search.

21. A program product having machine readable code for performing speech recognition, the program code, when executed, causing a machine to perform the following steps:

obtaining a list of target speech element sequences each containing at least one speech element;

for each target speech element sequence obtaining a forward sequence extension model and a backward sequence extension model;

spotting at least one spotted target speech element sequence in a set of acoustic observations by matching the sequence of speech element models against the set of acoustic observations;

obtaining from the set of acoustic observations the set of acoustic observations preceding the said at least one spotted target speech element sequence and the set of acoustic observations following the said at least one spotted target speech element sequence;

obtaining at least one hypothesis of a longer speech element sequence containing the said at least one spotted speech element sequence as a proper subsequence in which said at least one longer speech element sequence is consistent with at least one of said forward sequence extension model and said backward sequence extension model for said at least one spotted speech element sequence; and

evaluating said at least one hypothesis of a longer speech element sequence based on the degree of acoustic

match between said longer speech element sequence and at least one of said set of acoustic observations preceding the said at least one spotted target speech element sequence and the set of acoustic observations following the said at least one spotted target speech element sequence.

22. A program product as in claim 21, the program code further causing a machine to perform the following steps:

spotting a plurality of spotted target speech element sequences in the set of acoustic observations;

determining, for each spotted speech element sequence and each hypothesized longer speech element sequence, the set of acoustic observations that correspond to the speech interval for said speech element sequence;

detecting when the set of acoustic observations for a first speech element sequence and the set of acoustic observations for a second speech element sequence correspond to adjacent speech intervals; and

creating a combined speech element sequence by concatenating said first speech element sequence and said second speech element sequence.

23. A program product as in claim 21, the program code further causing a machine to perform the following steps:

obtaining from the set of acoustic observations the set of acoustic observations preceding the said at least one combined speech element sequence and the set of acoustic observations following the said at least one combined speech element sequence;

obtaining at least one hypothesis of a longer speech element sequence containing the said at least one combined speech element sequence as a proper subsequence in which said at least one longer speech element sequence is consistent with at least one of said forward sequence extension model of the spotted target speech element sequence contained in said second speech element sequence and said backward sequence

extension model for the spotted target speech element sequence contained in said first speech element sequence; and

evaluating said at least one hypothesis of a longer speech element sequence based on the degree of acoustic match between said longer speech element sequence and at least one of said set of acoustic observations preceding the said at least one combined speech element sequence and the set of acoustic observations following the said at least one combined speech element sequence.

24. A program product as in claim 21, the program code further causing a machine to perform the following steps:

repeating said processes of obtaining at least one hypothesis of a longer speech element sequence, and said evaluating said at least one hypothesis, and said determining of said sets of corresponding acoustic observations, until there is at least one pair of a first speech element sequence and a second element sequence for which it is detected that said first

speech element sequence and said second element sequence correspond to adjacent speech intervals;

creating said combined speech element sequence; and

repeating said processes of obtaining and evaluating said longer speech element sequences and of creating said combined speech element sequences until there is at least one hypothesized speech element sequence that corresponds to the complete set of acoustic observations.

25. A speech recognition method, comprising:

receiving a set of acoustic observations, and performing a speech recognition on the set of acoustic observations;

at the same time the speech recognition is being performed, determining whether or not an n-gram of speech elements occurs in the set of acoustic observations, wherein n is an integer greater than or equal to one;

if the determination is that an n-gram occurs, then performing at least one of a backward search and a forward

search using a continuation tree that represents allowable continuations in a grammar that may precede or follow the spotted n-gram; and

determining a best matching path in the continuation tree with respect to the set of acoustic observations.